Author Names & Affiliations

- Stuart Anderson - LIGO Laboratory, California Institute of Technology
- Patrick Brady - LIGO Scientific Collaboration, University of Wisconsin-Milwaukee
- Peter Couvares - LIGO Laboratory, California Institute of Technology
- Tom Downes - LIGO Scientific Collaboration, University of Wisconsin-Milwaukee
- Ryan Fisher - LIGO Scientific Collaboration, Syracuse University
- Albert Lazzarini - LIGO Laboratory, California Institute of Technology
- David Reitze - LIGO Laboratory, California Institute of Technology

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Physics, Astrophysics, Gravitational-Wave Astrophysics; Computer Science

Title of Submission

LIGO Response to NSF CI 2030 RFI

Abstract (maximum ~200 words).

The Laser Interferometer Gravitational-wave Observatory (LIGO) has detected gravitational waves and is eagerly awaiting an international network of detectors and the transition to an era of regular astrophysical observations that requires significant investment in Cyberinfrastructure.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The direct observation of gravitational waves from colliding black holes has opened a new window on the Universe. With just two confirmed detections in hand, the field of gravitational-wave astronomy is already living up to its promise to deliver new information about black holes, how they form, and how the Universe evolved. This is just the beginning. Over the next five years, LIGO and its international partners expect to improve the sensitivity of the current instruments to bring new observations at rates as high as once per day.

Looking to the future, the era CI 2030, the need for enhanced Cyberinfrastructure will be driven by four emerging opportunities (challenges): 1. Improvements in the low-frequency sensitivity of the detectors brings with it an increase in the number matched-filter

waveforms that must be compared to the data. Together with expansion of the search parameter space to include precession and ellipticity in the binary merger, this will drive a substantial increase in raw computing cycles needed to identify the signals in the data stream. 2. Increases in the number of gravitational-wave detectors at geographically distributed locations will allow better localization of the sources and enhanced tests of gravitational theories by adding polarization information. This in turn drives the need for robust and synchronized, low-latency data analysis both at the detector sites and in aggregate at central processing locations. 3. Increases in the number of international collaborations which work together on the global network of gravitational-wave detectors will increase the number of people who need to interact efficiently at all times of the day and night. 4. The number of signals will increase to many per day bringing a wealth of new information individually and in aggregate. It is also hoped to identify continuous (coherent and stochastic) signals and unmodeled transients that have not yet been found. The complexity and requirements on data processing will be substantially increased as a result.

The primary challenges related to Cyberinfrastructure are therefore:

* Growing computational scale (e.g., searching for binary coalescing systems at lower frequencies with possible orbital precession). Increasingly distributed, cross-organizational, and dynamic computing -- within and between GW collaborations and external computing services -- that requires flexible data and identity management.

* Lower-latency data analysis (e.g., detecting merging binary systems before they coalesce to point telescopes ahead of time) using data collected at multiple sites around the world.

* Increasingly complex integration and support of multiple loosely-coupled CI components to provide production services, rather than single, centralized stand-alone solutions. For example: integrating Shibboleth, Grouper, Kerberos, Sympa, LDAP, etc. into a unified Identity

* Access and Management solution; and integrating HTCondor, GlideinWMS, Globus, Xrootd, CVMFS, etc. into a unified grid computing infrastructure.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

1. Flexible data management for opportunistic computing. Advanced data transfer, data security, and data management capabilities to enable the flexible and seamless sharing of both data within GW collaborations and between collaborations, as well as the use of campus, regional, national, and commercial shared computing resources outside the collaborations. These capabilities should be well-integrated with existing tools and services widely used by scientists (for example workflow management systems, batch scheduling systems, storage management systems, and networking services) and should be well-supported as production resources. Potential challenges to leveraging opportunistic computing include data provenance, data integrity, data privacy, data replication, bandwidth-aware and storage-aware data transfer, accounting, and co-scheduling of data transfer, data storage, and processing. This is as much an integration, support, and sustainability challenge as it is a technical challenge.

2. Collaboration as a service. LIGO is developing, integrating, and supporting a wide array of coordinated collaboration tools that may not be inherently LIGO-specific, including identity and access management services (for onboarding/offboarding, group management, and access control), wikis, mailing-lists, voting systems, document management, source code version control, teleconferencing, chat, etc. It may be more efficient for this suite of standard collaboration tools to be provided as an NSF service to large science projects of comparable scale.

3. Open data and science gateway/hub. As the field of GW astronomy transitions from a phase of initial discoveries to regular astronomical observations it seems likely that standard science gateways will be developed for well-understood astrophysical sources and data analysis techniques, especially for data once they are publicly released. CI to support this would be helpful. For example, a national open data hub as a service may help maximize the science return.

4. Tools to support optimization for massively, multi-core hardware: The rapid increase in number of threads of execution per device is driving changes in the implementation of algorithms for scientific computing. To derive maximum benefit from this increase in computing

power without having to port each algorithm separately to new generations of hardware, the development of new languages and compilers that can automatically leverage these massively, multi-core environments is needed. There already exist a number of efforts to develop tools of these kinds, but they are not broadly adopted at this time for production scientific applications. An essential parallel effort will be the long-term support of these tools and languages in addition to education and training to use them. At the moment, innovation in hardware is moving faster than our abstractions are keeping up. This is a fundamental area in which the NSF CI 2030 program could have substantial impact for the future of gravitational waves.

5. Tools to support code optimization for high-throughput computing environments. Frameworks, tools, services, and dissemination of best practices for the development, building, packaging, and deployment of hardware-optimized code on distributed resources. LIGO seeks to efficiently use the most cost-effective computing resources available, and to enable elasticity of the supply of computing resources available for GW data analysis to meet uneven or unpredictable demand. Challenges include providing heterogenous build environments and/or cross-compiling, code and hardware identification and matchmaking, and software dependency management (e.g., via matchmaking, packaging, and/or containerization). This should support multiple programming and scripting languages, CPU, GPU, and co-processor platforms, open source and commercial development tools (e.g., icc/MKL, CUDA, etc.).

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

1. Sustaining cyberinfrastructure: Cyberinfrastructure, including software packages developed both inside and outside the collaboration, are essential to the success of LIGO. It is critically important that NSF plan for long-term, sustainable funding of this CI that matches the lifetime of the NSF large facilities and research programs. As identified in the first NSF Large Facilities Cyberinfrastructure Workshop in 2015, there remains a matchmaking problem between the developers of NSF-funded CI and the potential scientific end-users. Making it easier to know what CI exists and how well supported it is may reduce the number of times researchers have to reinvent the wheel, or reduce the number of times researchers invest in short-lived externally-developed tools when they should be reinventing their own.

2. Support for integration of emerging technologies for scientific discovery: Funding long-term projects to host and support multiple coordinated CI components as a production service, providing flexible integration services to facilitate adoption, and consulting to enable a migration to self-hosted services when appropriate.

3. Sustaining a well-trained workforce: Moreover, the integration, maintenance and operations of this CI for large scientific collaborations such as LIGO requires data scientists with skills in both the domain and in cyberinfrastructure. At the present time, the need for support is outpacing the available funding for people. Solicitations could increasingly account for the need for retaining and supporting well-trained scientists who specialize in infrastructure to reduce the cost of retraining repeatedly, and minimize the loss of expertise to industry.

4. International obstacles to large-scale collaboration: Establishing international agreements to facilitate the use of federated identities across national boundaries, including issues of privacy rules and regulations. In particular, the complex landscape of international laws relating to data privacy are difficult to address at the Collaboration or institutional level. This is an area where the NSF could provide leadership and advocacy to ensure that scientific advancement that relies on international collaboration can be efficiently carried out.

Consent Statement